Analysing the Effects of Routing Centralization on BGP Convergence Time

Pavlos Sermpezis
FORTH, Greece
sermpezis@ics.forth.gr

Xenofontas Dimitropoulos
FORTH, Greece
fontas@ics.forth.gr

arXiv:1605.08864v1 [cs.NI] 28 May 2016

1. INTRODUCTION

Software-defined networking (SDN) has improved the routing functionality in networks like data centers or WANs. Recently, several studies proposed to apply the SDN principles in the Internet's inter-domain routing as well [1–5]. This could offer new routing opportunities [1, 2] and improve the performance of BGP [3–5], which can take minutes to converge to routing changes [6–8].

Previous works have demonstrated that centralization can benefit the functionality of BGP, and improve its slow convergence that causes severe packet losses [6] and performance degradation [7]. However, due to (a) the fact that previous works mainly focus on system design aspects, and (b) the lack of real deployments, it is not clearly understood yet *to what extent* inter-domain SDN can improve performance.

To this end, in this work, we make the first effort towards *analytically* studying the effects of routing centralization on the performance of inter-domain routing, and, in particular, the convergence time of BGP. Specifically, we propose a Markovian model for inter-domain networks, where a subset of nodes (*domains*) coordinate to centralize their inter-domain routing. We then derive analytic results that quantify the BGP convergence time under various network settings (like, SDN penetration, topology, BGP configuration, etc.). Our analysis and results facilitate the performance evaluation of inter-domain SDN networks, which have been studied (till now) only through simulations/emulations that are known to suffer from high time/resource requirements and limited scalability.

2. MODEL

Network Model. We consider a network (e.g., the Internet) composed of N *domains* or *autonomous systems* (ASes). We represent each AS as a single node, i.e., a single BGP router (similarly to [3]). Such an abstraction allows to hide the details of the internal structure of ASes, and focus on inter-domain routing.

We assume that $k \in [1, N]$ ASes cooperate in order to centralize their inter-domain routing: there exists a *multi-domain SDN controller*, which is connected to the BGP routers of these k ASes[1]. In the remainder, we refer to the set of the k ASes, as the *SDN cluster*.

BGP Updates. As in the Internet, ASes use BGP to exchange information and establish routing paths. When a BGP edge router of an AS receives a BGP update, it (i) calculates the updates (if any) for its BGP routing table, (ii) sends updates to the other BGP edge routers within the same AS (e.g., with iBGP), and (iii) sends updates to the BGP routers of the neighboring ASes. The time needed for this process may vary a lot among different connections since it depends on a number of factors, like the employed technology (hardware/software), routers' configuration (e.g., MRAI timers), intra-domain network, etc. To this end, in order to be able to analytically study the BGP updates dissemination (given the uncertainty and complexity), we model the time between the reception and forwarding of a BGP update in a probabilistic way.

Assumption 1. *The time between the reception of a BGP update in an AS/router and its forwarding to a neighbor AS/router, is exponentially distributed with rate λ.*

Despite the simplicity of the above assumption, our results can capture the behavior of real/emulated networks (see Section 4).

Inter-domain SDN routing. Each AS belonging to the SDN cluster informs the SDN controller upon the reception of a BGP update. The SDN controller, which is aware of the topology of the SDN cluster (neighbors, policies, paths, etc.), calculates the changes in the routing paths and installs the updated routes in each router/AS belonging to the SDN cluster. ASes react to updates from the SDN controller, as in regular BGP updates, and, thus, forward them to their (non SDN) neighbors.

Let T_{sdn} be the time needed for an AS to inform the SDN controller and the controller to install the updated routes in every AS in the SDN cluster. This time can be expected to be in the order of few seconds [3], and much lower than the BGP updating process (cf., the default value for MRAI timers in Cisco routers is $30sec.$), thus, for simplicity, we assume here that $T_{sdn} = 0$.

3. ANALYSIS: BGP CONVERGENCE TIME

Let us assume a routing change, e.g., an announcement of a new prefix by an AS, in the network at time $t_0 = 0$. Our goal is to calculate the *BGP convergence time*, i.e., the time needed till all ASes/routers in the network have the final (i.e., shortest, conforming to policies) BGP routes for this prefix.

To this end, using Assumption 1, we can model the dissemination of the BGP updates in the network with the Markov Chain (MC) of Fig. 1(a), where each state corresponds to the number of ASes/routers that have the final BGP routes. At time $t_0 = 0$ the system is at state 0, while the state N denotes the BGP convergence. When an AS in the SDN cluster receives the BGP update, all the nodes in the SDN cluster are informed (through the controller); thus, we have a transition, e.g., from state i to state $k + i$. The transition rates in the MC, as we discuss in detail later, depend on the network topology.

The Markov Chain of Fig. 1(a) is transient, and the BGP convergence time is the time needed to move from state 0 to state N.

For notation brevity, in the remainder we use the MC of Fig. 1(b),

[1]This system abstraction can capture the main functionality of most of the previously proposed approaches.

Copyright is held by author/owner(s).

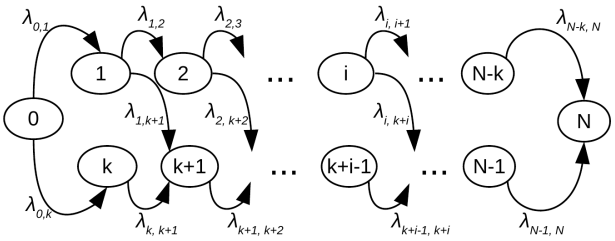

(a) Markov Chain (number of nodes)

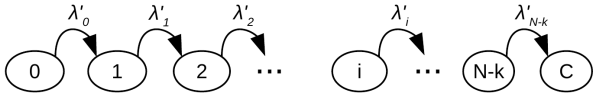

(b) Markov Chain (number of transitions, or *steps*)

Figure 1: Markov Chains where the states correspond to (a) the number of nodes that have updated BGP routes, and (b) the number of transitions, or *steps*, of the BGP update dissemination process.

which is equivalent to the MC of Fig. 1(a). Here, the states represent the *number of transitions* in the MC of Fig. 1(a). For example, the state/step 1 corresponds to the state 1 or k of the MC of Fig. 1(a), while the state/step i corresponds to the state i or $k+i-1$ in the MC of Fig. 1(a). The states 0 are equivalent in both MCs, while the state/step C denotes the BGP convergence, and, thus, corresponds to the state N in the MC of Fig. 1(a).

If we denote with x the step at which -for the first time- an AS in the SDN cluster receives the BGP update, then the transitions rates $\lambda_i^{'}$ in the MC of Fig. 1(b) are given by

$$\lambda_i^{'} = \begin{cases} \lambda_{i,i+1} + \lambda_{i,i+k} & , i \leq x \\ \lambda_{k+i-1,k+i} & , i > x \end{cases} \quad (1)$$

We now proceed to calculate the rates $\lambda_i^{'}$. The ASes that have received the BGP updates, will then send BGP updates to some of their neighboring ASes, according to their routing policies. We refer to a neighbor to which the update will be forwarded as a *bgp-eligible* neighbor.

Definition 1. *We define as the* `bgp-degree` *at step i, $D(i)$, the number of the ASes that are bgp-eligible neighbors with any of the ASes that have received the BGP updates at step i.*

Although an AS might receive the same BGP update from more than one neighbors, the final BGP route will correspond to only one of the received updates (i.e., shortest path). Hence, to calculate the transition rate $\lambda_i^{'}$, we take into account only one BGP connection (corresponding to the shortest path) per bgp-eligible neighbor. Since the BGP update times are exponentially distributed with rate λ (see, Assumption 1), it follows that $\lambda_i^{'}$ will be given by[2]

$$\lambda_i^{'} = \lambda \cdot D(i) \quad (2)$$

Knowing the rates $\lambda_i^{'}$, we can calculate the transition delays in each step. Adding the delays in each step, we can derive Theorem 1, which gives the BGP convergence time, i.e., the time to move from state 0 to state C.

[2]The transition time is the minimum of $D(i)$ i.i.d. exponentially distributed times.

Theorem 1. *The expectation of the BGP convergence time T in a hybrid SDN/BGP inter-domain topology is given by*

$$E[T] = \frac{1}{\lambda} \cdot \sum_{x=0}^{N-k} \sum_{i=1}^{N-k} \frac{1}{D(i|x)} \cdot P_{sdn}(x) \quad (3)$$

where $D(i|x)$ is the `bgp-degree` *of the network at step i given that the SDN cluster receives the update at step x, and $P_{sdn}(x)$ is the probability that the SDN cluster receives the update at step x.*

Proof. To calculate $E[T]$, we first apply the conditional expectation

$$E[T] = \sum_{x=0}^{N-k} E[T|x] \cdot P_{sdn}(x) \quad (4)$$

where $E[T|x]$ denotes the expected convergence time, given that the SDN cluster receives the update at step x. Under this condition, the bgp-degrees at each step are $D(i|x)$, $i \in [1, N-k]$. Hence, taking also into account Eq. (2), it follows that the transition delay from a step i to a step $i+1$, $T_{i,i+1}$, is exponentially distributed with rate $\lambda_i^{'} = \lambda \cdot D(i|x)$, and its expectation is given by

$$E[T_{i,i+1}|x] = \frac{1}{\lambda \cdot D(i|x)} \quad (5)$$

Remark: The state/step 0 does not correspond to a real state of the system; it is only used for presentation purposes. Thus, we set $T_{0,1} = 0$.

As mentioned earlier, the BGP convergence delay is the time needed to move from step 0 to step C, and thus it is given by the sum of the transition delays of all the intermediate steps, i.e.,

$$\begin{aligned} E[T|x] &= E\left[\sum_{i=1}^{N-k} T_{i,i+1}|x\right] \\ &= \sum_{i=1}^{N-k} E[T_{i,i+1}|x] \\ &= \sum_{i=1}^{N-k} \frac{1}{\lambda \cdot D(i|x)} \end{aligned} \quad (6)$$

Now, the expression of Eq. (3) follows by substituting Eq. (6) to Eq. (4). □

In the following sections we calculate the quantities $D(i|x)$ and $P_{sdn}(x)$ for important network topologies.

3.1 Full-Mesh Network Topology

We first consider a basic topology: a full-mesh network, where every AS-pair is connected.

Theorem 2. *The probability that the SDN cluster receives the update at step x is given by*

$$P_{sdn}(x) = \frac{k}{N-x} \cdot \prod_{j=0}^{x-1} \left(1 - \frac{k}{N-j}\right) \quad (7)$$

Proof. The SDN cluster comprises k (out of the total N) ASes. Since we consider that the prefix announcement is made by a (random) AS in the network, the probability that the announcing AS is in the SDN cluster (and thus $x = 0$) is

$$P_{sdn}(0) = \frac{k}{N} \quad (8)$$

If the announcing AS is not in the SDN cluster, then $x > 0$, and thus

$$P_{sdn}(x > 0) = 1 - P_{sdn}(0) = 1 - \frac{k}{N} \tag{9}$$

The probability $P_{sdn}(1)$ is given by

$$P_{sdn}(1) = P_{sdn}(1|x > 0) \cdot P_{sdn}(x > 0) \tag{10}$$

where $P_{sdn}(1|x > 0)$ denotes the probability that the SDN cluster receives the BGP update at step 1, given that it has not received it before. If $x > 0$, then at step 1 the remaining ASes without the update are $N - 1$, of which k belong to the SDN cluster. Since the BGP update processes are distributed with the same rate λ, the probability that the next AS to get the update belongs to the SDN cluster is $\frac{k}{N-1}$. Therefore, and taking into account Eq. (9), it holds that

$$P_{sdn}(1) = P_{sdn}(1|x > 0) \cdot P_{sdn}(x > 0) = \frac{k}{N-1} \cdot \left(1 - \frac{k}{N}\right) \tag{11}$$

and, respectively,

$$\begin{aligned} P_{sdn}(x > 1) &= (1 - P_{sdn}(1|x > 0)) \cdot P_{sdn}(x > 0) \\ &= \left(1 - \frac{k}{N-1}\right) \cdot \left(1 - \frac{k}{N}\right) \end{aligned} \tag{12}$$

Proceeding similarly for the next steps $i = 2, ..., N - k$, it can be shown that

$$P_{sdn}(i|x > i - 1) = \frac{k}{N - i} \tag{13}$$

and

$$\begin{aligned} P_{sdn}(x > i - 1) &= (1 - P_{sdn}(i - 1|x > i - 2)) \cdot ... \cdot P_{sdn}(x > 0) \\ &= \left(1 - \frac{k}{N - (i-1)}\right) \cdot ... \cdot \left(1 - \frac{k}{N}\right) \\ &= \prod_{j=0}^{i-1} \left(1 - \frac{k}{N-j}\right) \end{aligned} \tag{14}$$

and, therefore,

$$\begin{aligned} P_{sdn}(i) &= P_{sdn}(i|x > i - 1) \cdot P_{sdn}(x > i - 1) \\ &= \frac{k}{N-i} \cdot \prod_{j=0}^{i-1} \left(1 - \frac{k}{N-j}\right) \end{aligned} \tag{15}$$

which is the expression of Eq. (7). □

Theorem 3 gives the bgp-degrees $D(i|x)$ in a mesh network as a function of $n(i|x)$, which is defined as the number of nodes with updated BGP information at step i, given that the SDN cluster received the update at step x

$$n(i|x) = \begin{cases} i & , i \leq x \\ i + k - 1 & , i > x \end{cases} \tag{16}$$

Theorem 3. *The* `bgp-degree` $D(i|x)$, $i \in [1, N - k], x \in [0, N - k]$, *in a full-mesh network topology is given by*

$$D(i|x) = N - n(i|x) \tag{17}$$

Proof. In a mesh network, since every AS-pair is directly connected, only the BGP messages sent by the announcing AS (i.e., shortest path) need to be considered. In step i, the announcing AS has $N - n(i|x)$ neighbors that have not received the BGP updates, and thus, it follows that $D(i|x) = N - n(i|x)$. □

3.2 Random Graph Network Topologies

In networks that are not full-meshes, ASes can be connected in many different ways and policies. Since it is not possible to study every single topology, we use two classes of random graphs to capture the effects of routing centralization in non full-mesh networks. In this first approach, we consider unconstrained routing policies.

3.2.1 Poisson (Erdos-Renyi) Graph

We first consider the case of a Poisson random graph, where a link between each AS-pair exists with probability p. Varying the value of p we can capture different levels of sparseness.

Using similar arguments as in the full-mesh case, it is easy to show that the probabilities $P_{sdn}(x)$ are given by Theorem 2. The *expected* bgp-degrees, which can be used (as an approximation) instead of $D(i|x)$ in Theorem 1, are given by the following Theorem.

Theorem 4. *The expectation of the* `bgp-degree` $D(i|x)$, $i \in [1, N - k], x \in [0, N - k]$, *in a Poisson graph network topology is*

$$E[D(i|x)] = (N - n(i|x)) \cdot \left(1 - (1 - p)^{n(i|x)}\right) \tag{18}$$

Proof. In a non full-mesh network, some ASes are not directly connected to the announcing AS. Thus, in the calculation of the $D(i|x)$ we need to consider the bgp-eligible neighbors of *all* the ASes that have received the update.

Let assume that we are at step i, and $n(i)$ nodes have received the BGP updates; we denote the set of these nodes as S_i. A node $m \notin S_i$ is a bgp-eligible neighbor with a node $j \in S_i$ with probability

$$P(m, j) = p \tag{19}$$

since every pair of nodes is connected with probability p (by the definition of a Poisson graph). The probability that m is *not* a bgp-eligible neighbor with *any* of the nodes $j \in S_i$, is given by

$$1 - P(m, S_i) = \prod_{j \in S_i} (1 - P(m, j)) = \prod_{j \in S_i} (1 - p) = (1 - p)^{n(i)} \tag{20}$$

since $|S_i| = n(i)$. It follows easily that the complementary event, i.e., m is a bgp-eligible neighbor with *any* of the nodes $j \in S_i$, happens with probability

$$P(m, S_i) = 1 - (1 - p)^{n(i)} \tag{21}$$

There are $N - n(i)$ ASes without the update, with each of them being a bgp-eligible neighbor with any of the nodes $j \in S_i$ with (equal) probability $P(m, S_i)$. Hence, the total number of bgp-eligible neighbors (or, as defined in Def. 1, the *bgp-degree* $D(i)$) is a binomially distributed random variable, whose expectation is given by

$$E[D(i)] = (N - n(i)) \cdot (1 - (1 - p)^{n(i)}) \tag{22}$$

□

Corollary 1. *Using the expectation of $D(i|x)$ in Theorem 1, underestimates the BGP convergence time, i.e.,*

$$E[T] \geq \frac{1}{\lambda} \cdot \sum_{x=0}^{N-k} \sum_{i=1}^{N-k} \frac{1}{E[D(i|x)]} \cdot P_{sdn}(x) \tag{23}$$

Proof. The bgp-degree $D(i|x)$ in non full-mesh networks, is a random variable that can take different values, depending on the BGP updates dissemination process (i.e., the exact set of nodes that have

received the BGP updates at step i, and their links to the rest of the nodes). Thus, we can write for the transition delay

$$E[T_{i,i+1}|x] = \sum_y \frac{1}{y} \cdot P\{D(i|x) = y\} = E\left[\frac{1}{D(i|x)}\right] \quad (24)$$

Calculating the exact value of the expectation $E\left[\frac{1}{D(i|x)}\right]$ is difficult, thus, we use a well known approximation in Eq. (24):

$$E[T_{i,i+1}|x] = E\left[\frac{1}{D(i|x)}\right] \approx \frac{1}{E[D(i|x)]} \quad (25)$$

where the calculation of $E[D(i|x)]$ is much easier (see, e.g., proof of Theorem 4). Then the BGP convergence delay is approximately given by

$$E[T] \approx \frac{1}{\lambda} \cdot \sum_{x=0}^{N-k} \sum_{i=1}^{N-k} \frac{1}{E[D(i|x)]} \cdot P_{sdn}(x) \quad (26)$$

However, applying Jensen's bound for the expectation of a convex function (here, $f(x) = \frac{1}{x}$) of a random variable (here, $D(i|x)$) on Eq. (24), gives

$$E[T_{i,i+1}|x] = E\left[\frac{1}{D(i|x)}\right] \geq \frac{1}{E[D(i|x)]} \quad (27)$$

which proves the Corollary. □

3.2.2 Arbitrary Degree Sequence Random Graph

The structure of networks where the degrees (i.e., the number of connections) of the ASes are largely heterogeneous, e.g., power-law graphs, can be better described with a Configuration-Model Random Graph (CM-RG) rather than a Poisson graph. In the CM-RG model, a random graph is created by connecting randomly the nodes (i.e., ASes), whose degrees are given [9]. Hence, we can use the CM-RG to model a network with *any arbitrary degree sequence* with mean value μ_d and variance σ_d^2 (and, $CV_d = \frac{\sigma_d}{\mu_d}$).

If the participation of an AS in the SDN cluster is independent of its degree, then the probabilities $P_{sdn}(x)$ are given by Theorem 2, and the bgp-degrees $D(i|x)$ are given by the following Result[3].

Result 1. *The expectation of the* `bgp-degree` $D(i|x), i \in [1, N-k], x \in [0, N-k]$, *in a CM-RG network topology is given by*

$$E[D(i|x)] = D(1|x) \cdot \prod_{j=1}^{i-1} A(j|x) + \sum_{j=1}^{i-1} (\mu_d(j|x) - 1) \cdot \prod_{m=j+1}^{i-1} A(m|x) \quad (28)$$

where

$$D(1|x) = \begin{cases} \mu_d & , x > 0 \\ (N-k) \cdot \mu_d \cdot \ln\left(\frac{N}{N-k}\right) & , x = 0 \end{cases} \quad (29)$$

$$\mu_d(j|x) = \mu_d \cdot \prod_{m=1}^{j-1} \left(1 - \frac{CV_d^2}{N - n(m|x) - 1}\right) \quad (30)$$

$$A(j|x) = 1 - \frac{\mu_d(j|x)}{N - n(j|x) - 1} \quad (31)$$

Proof. The main difference with the Poisson case is that in the CM-RG case, it is more probable that the ASes with the higher degrees will receive the BGP updates faster. For instance, let us assume that the announcing AS, e.g., AS-1, does not belong to the SDN cluster. If we denote with d_1, d_2, and d_3 the degrees of AS-1, AS-2 and AS-3 (where AS-2 and AS-3, have not received yet the BGP update), a property of a CM-RG says that AS-1 is directly connected with AS-2 and AS-3 with probabilities

$$P(1,2) = c \cdot d_1 \cdot d_2 \quad and \quad P(1,3) = c \cdot d_1 \cdot d_3 \quad (32)$$

respectively, where c a normalizing constant. In other words, the AS with with the higher degree has a higher probability to be directly connected to AS-1. Consequently, ASes with higher degrees have a higher probability to get the BGP update faster.

Now, let us first derive Eq. (29). If $x > 0$, the announcing AS does not belong to the SDN cluster ($n(1|x > 0) = 1$), and thus the bgp-degree will be equal to the degree of the announcing AS. Since the average degree of a node is μ_d, it follows easily that expectation of the bgp-degree in this case, is given by

$$E[D(1|x > 0)] = \mu_d \quad (33)$$

If $x = 0$, the announcing AS belongs to the SDN cluster, and, thus, all the k nodes in the SDN cluster have the BGP updates. In this case, the bgp-degree is the number of all bgp-eligible neighbors of these k nodes. Let us denote with S_1 the set of nodes in the SDN cluster. Since the fact that a node belongs to the SDN cluster and its degree are independent, the probability that an edge coming out from a node $m \notin S_1$ is connected to a node $j \in S_1$, is equal to $\frac{k}{N}$. Hence, a node $m \notin S_1$, with degree d_m, is not connected to any of the k nodes in the SDN cluster with probability

$$1 - P(m, S_1) = \left(1 - \frac{k}{N}\right)^{d_m} \quad (34)$$

and, respectively

$$P(m, S_1) = 1 - \left(1 - \frac{k}{N}\right)^{d_m} \quad (35)$$

The above equation holds $\forall m \notin S_1$; the degrees d_m can have different values for each m. Since, there are $N-k$ nodes that do not belong to the SDN cluster (and, do not have the BGP updated route), the expected bgp-degree is given by

$$E[D(1|0)] = (N-k) \cdot E\left[1 - \left(1 - \frac{k}{N}\right)^d\right] \quad (36)$$

where the expectation is taken over d, i.e., over all the degrees $d_m, m \notin S_1$. To calculate this expectation is difficult, thus we approximate it using a Taylor series approximation, i.e,

$$\begin{aligned} E\left[1 - \left(1 - \frac{k}{N}\right)^d\right] &= 1 - E\left[\left(1 - \frac{k}{N}\right)^d\right] \\ &\approx 1 - \left(1 + E[d] \cdot \ln\left(1 - \frac{k}{N}\right)\right) \\ &= -E[d] \cdot \ln\left(1 - \frac{k}{N}\right) \\ &= E[d] \cdot \ln\left(\frac{N}{N-k}\right) \\ &= \mu_d \cdot \ln\left(\frac{N}{N-k}\right) \end{aligned} \quad (37)$$

which completes the derivation of Eq. (29).

To compute the bgp-degrees $D(i|x)$ of the steps $i = 2, ..., N-k$, we follow a methodology similar to [10]. Let $D(i-1)$ be the bgp-degree at step $i-1$ and $\mu_d(i-1)$ the average degree of the nodes that have not received the BGP update by step $i-1$ (i.e., the nodes

[3]We use the notation "Result", instead of "Theorem", because the provided expression is an approximation.

$m, m \notin S_{i-1}$). The average degree $\mu_d(i-1)$ is not equal to μ_d, in general; this is due to the fact that nodes with higher degrees receive faster the BGP updates and thus the remaining nodes are nodes with lower degrees (for a more detailed argumentation see [10]). Following similar arguments as in [10], it can be shown that the average degrees $\mu_d(i)$ are approximately given by Eq. (30).

We calculate the bgp-degree of the step i, based on the quantities $D(i-1)$ and $\mu_d(i-1)$, as follows

$$D(i) = D(i-1) - 1 + \mu_d(i-1) \cdot \left(1 - \frac{D(i-1)}{N - n(i-1)}\right) \quad (38)$$

The term $D(i-1)-1$ is the bgp-degree of the previous step minus 1, which denotes the i^{th} node that received the BGP update (i.e., at the transition between step $i-1$ and i). To this quantity, we need to add the number of nodes that are bgp-eligible neighbors of the i^{th} node, but are not bgp-eligible neighbors of any of the nodes $\in S_{i-1}$. The total nodes without the updated BGP routes at step $i-1$ are $N - n(i-1)$; $D(i-1)$ of these nodes are bgp-eligible neighbors of at least one node $\in S_{i-1}$. Hence, the probability that a bgp-eligible neighbor of the i^{th} node is not a bgp-eligible neighbor of any of the nodes $\in S_{i-1}$, is given by $\left(1 - \frac{D(i-1)}{N-n(i-1)}\right)$. Since, the i^{th} node has (on average) $\mu_d(i-1)$ bgp-eligible neighbors, it follows that the number we need to add to the quantity $(D(i-1)-1)$ is

$$\mu_d(i-1) \cdot \left(1 - \frac{D(i-1)}{N - n(i-1)}\right) \quad (39)$$

Finally, calculating recursively Eq. (38), after some algebraic manipulations, we derive Eq. (28).

□

4. VALIDATION AND DISCUSSION

We built a simulator, conforming to our model (i.e., Assumption 1, $T_{sdn} = 0$, etc.). In Fig. 2 we compare our theoretical results against simulations (averages over 200 runs). The accuracy is high for the Full-mesh and Poisson graph cases (Fig. 2(a)). In the CM-RG case (Fig. 2(b)) the two curves are similar, and the error of our expression (which is an approximation) is always less than 18%.

Moreover, an initial comparison with the emulation results of [3], where a *real BGP software router* is employed, shows that our theory (despite the assumptions and model simplicity) is in agreement with their observations: e.g., the convergence time (i) decreases faster after $\frac{k}{N} > 50\%$, (ii) has small differences among different topologies when centralization exists, etc.

This highlights one of the contributions of our work: with our results, we can quickly evaluate the effects of routing centralization. The provided expressions are simple (need only to know a few parameters: N and k, and -if needed- p, or μ_d and CV_d) and easy/fast to compute, whereas emulations are time/resource demanding and have limited scalability. Hence, one could use our work to obtain initial insights, and then, e.g., proceed to a few targeted emulations for a more fine-grained or system-specific investigation.

A second contribution is that our methodology and results can be used as the building blocks to model and analyse more complex settings; i.e., model different parts of a large network as full-mesh/Poisson/CM-RG sub-networks, use our expressions for each sub-network, and synthesize them to compute the overall network performance. We consider such an example in Section 5, where we model the core of the Internet, considering different classes of tier-1 and tier-2 ISPs, while taking into account their routing policies as well.

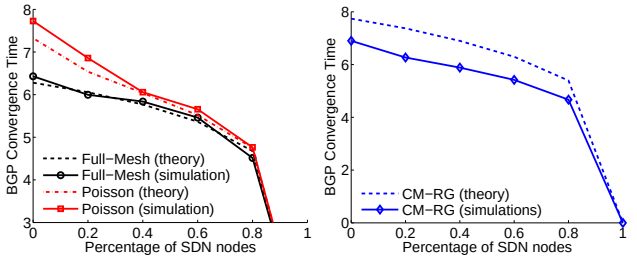

(a) (b)

Figure 2: BGP convergence time vs. SDN penetration (i.e, $\frac{k}{N}$) in networks with $N = 300$. (a) Full-mesh and Poisson graph ($p = \frac{1}{60}$) topologies; and (b) CM-RG topology with degree sequence d_i, $d_i \in [5, 200]$ and power-law distributed with exponent 2.

Finally, we believe that this first analytic study, offers a new useful approach for investigating inter-domain SDN, and can be the basis for analysing further aspects of this field. In particular, we plan to extend our methodology, and consider more settings (e.g., BGP update times), performance metrics, and applications.

5. CASE STUDY: BGP CONVERGENCE AT THE INTERNET CORE NETWORK

In this section, we use our results to study the effects of routing centralization in the core of the Internet. We show that our analysis can provide useful insights related to the evolution of the Internet, and answer questions like "*Internet flattening* or *routing centralization* can improve more the BGP convergence?".

We first need to model the core of the Internet. We consider the following setting, which captures the structure in the core of the Internet in an (a) abstract (and thus analytically tractable), (b) generic (and thus easily adapted to network characteristics and changes), and (c) realistic (since it takes into account routing policies, e.g., customer to provider, peering links, valley-free model) way.

Specifically, we assume that the Internet core network is composed of N_1 tier-1 ISPs, and N_2 tier-2 ISPs. Let tier-1 ISPs to peer with each other with probability p_{11} (Poisson graph), while tier-2 ISP peer with probability p_{22}. Also, tier-1 ISPs offer transit to tier-2 ISPs; each tier-2 ISP might have one or more transit providers. Let p_{12} be the probability that a tier-1 and a tier-2 ISPs are connected.

Consider a prefix announced by AS-x, a tier-2 ISP. Due to the routing policies (peering/transit), the BGP updates are forwarded as follows. The tier-2 *peers* of AS-x receive the BGP updates directly from AS-x. The same holds also for the tier-1 ISPs that are connected to AS-x. The remaining tier-1 ISPs receive the updates from other tier-1 ISPs, and the remaining tier-2 ISPs from their tier-1 providers. Fig. 3 demonstrates such an example topology and the respective eligible paths of BGP updates.

Two recent trends, related to the evolution of the Internet, have been recently observed and/or proposed: (a) Internet flattening [11], i.e., more peer connections appear. In the above model, higher values of p_{22} correspond to more flattening. (b) Inter-domain SDN, where some ISPs centralize their routing. Network economics and feasibility studies indicate that routing centralization is more probable to start from larger ISPs [1]; hence, we assume that $k_1 \in [1, N_1]$ tier-1 ISPs belong to an SDN cluster.

Using the results of the previous sections, and making similar analytic arguments, in the remainder, we derive expressions for the BGP convergence time as a function of the network parameters $(N_1, k_1, N_2, p_{11}, p_{22}, p_{12})$. Therefore, it becomes straightforward

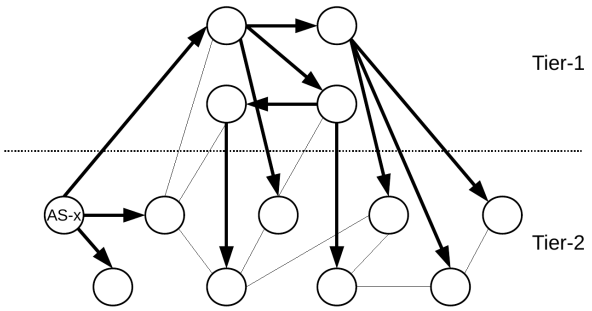

Figure 3: Internet core topology example. Lines/arrows denote connections between ASes. Arrows denote the eligible paths of BGP updates.

to compare the two approaches, Internet flattening vs. routing centralization, by evaluating the expressions for different values of p_{22} and k_1.

Let us use the notation

$$G_{fm}(N,k,\lambda) \quad \text{and} \quad G_p(N,k,p,\lambda) \tag{40}$$

for full-mesh and Poisson graph networks (with the given parameters), respectively.

The BGP convergence delay is given by the following Result.

Result 2. *The BGP convergence delay in the core of the Internet is given by*

$$E[T] = max\left\{E[T_{peering}], E[T_{transit}]\right\} \tag{41}$$

where

$$E[T_{peering}] = E[T|G_{fm}\left(N_2 \cdot p_{22}, 1, \lambda\right)] \tag{42}$$

$$E[T_{transit}] = E[T_{x,tier-1}] + E[T_{tier-1}] + E[T_{tier-1,tier-2}] \tag{43}$$

and

$$E[T_{x,tier-1}] = \frac{1}{p_{12} \cdot N_1} \tag{44}$$

$$E[T_{tier-1}] \approx E[T|G_p\left(N_1, k_1, p_{11}, \lambda\right)] \tag{45}$$

$$E[T_{tier-1,tier-2}] \approx E[T|G_{fm}\left(N_2 \cdot (1-p_{22}), 1, N_1 \cdot p_{12} \cdot \lambda\right)] \tag{46}$$

Proof. $E[T_{peering}]$ is the time till every peer of AS-x (i.e., direct connections to tier-2 ISPs) to receive the updates, and $E[T_{transit}]$ the respective time for the non-peering ASes (which get the updates through transit, i.e., the tier-1 ISPs).

AS-x has (on average) $N_2 \cdot p_{22}$ peers, which are directly connected to it. Hence, for the BGP update dissemination process we can use the full-mesh network model with $k = 1$ (i.e., without centralization), since we consider centralization only for tier-1 ISPs. As a result, it holds that

$$E[T_{peering}] = E[T|G_{fm}\left(N_2 \cdot p_{22}, 1, \lambda\right)] \tag{47}$$

To calculate $E[T_{transit}]$, we split it into three parts, i.e.,

$$E[T_{transit}] = E[T_{x,tier-1}] + E[T_{tier-1}] + E[T_{tier-1,tier-2}] \tag{48}$$

where (i) $T_{x,tier-1}$ is the time till the first tier-1 AS (a transit provider of AS-x) receives the BGP update, (ii) T_{tier-1} is the time, after $T_{x,tier-1}$, needed for every tier-1 ISP to get the BGP update, and (iii) $T_{tier-1,tier-2}$ is the time, after T_{tier-1}, needed for the tier-1 ASes to send the BGP updates to their customers (i.e., tier-2 ASes).

Using the Markovian properties of BGP update times and making similar arguments as before, we can show that

$$E[T_{x,tier-1}] = \frac{1}{p_{12} \cdot N_1} \tag{49}$$

and

$$E[T_{tier-1}] \approx E[T|G_p\left(N_1, k_1, p_{11}, \lambda\right)] \tag{50}$$

Also, it holds that

$$E[T_{tier-1,tier-2}] \approx E[T|G_{fm}\left(N_2 \cdot (1-p_{22}), 1, N_1 \cdot p_{12} \cdot \lambda\right)] \tag{51}$$

since there exist $N_2 \cdot (1-p_{22})$ tier-2 ASes that are not peers with AS-x (and thus will receive the update from tier-1 ASes), and each of them is (on average) connected to $N_1 \cdot p_{12}$ tier-1 ISPs (and can receive updates from any of them). □

6. REFERENCES

[1] Kotronis, V. and Kloeti, R. and Rost, M. and Georgopoulos, P. and Ager, B. and Schmidt, S. and Dimitropoulos, X., "Stitching inter-domain paths over IXPs," in *ACM SOSR*, 2016.

[2] A. Gupta, L. Vanbever, M. Shahbaz, S. Donovan, B. Schlinker, N. Feamster, J. Rexford, S. Shenker, R. Clark, and E. Katz-Bassett, "SDX: A software defined internet exchange," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 4, pp. 551–562, 2014.

[3] V. Kotronis, A. Gämperli, and X. Dimitropoulos, "Routing centralization across domains via sdn: A model and emulation framework for BGP evolution," *Computer Networks*, pp. –, 2015.

[4] C. Rothenberg, M. Nascimento, M. Salvador, C. Corrêa, S. C. de Lucena, and R. Raszuk, "Revisiting routing control platforms with the eyes and muscles of software-defined networking," in *Proc. ACM HotSDN*, 2012.

[5] M. Caesar, D. Caldwell, N. Feamster, J. Rexford, A. Shaikh, and J. van der Merwe, "Design and implementation of a routing control platform," in *Proc. NSDI*, 2005.

[6] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian, "Delayed Internet routing convergence," *ACM SIGCOMM Computer Communication Review*, vol. 30, no. 4, pp. 175–187, 2000.

[7] N. Kushman, S. Kandula, and D. Katabi, "Can you hear me now?!: it must be BGP," *ACM SIGCOMM Computer Communication Review*, vol. 37, no. 2, pp. 75–84, 2007.

[8] R. Oliveira, B. Zhang, D. Pei, and L. Zhang, "Quantifying path exploration in the Internet," *Networking, IEEE/ACM Transactions on*, vol. 17, no. 2, pp. 445–458, 2009.

[9] M. Newman, *Networks: An Introduction*. Oxford University Press, Inc., 2010.

[10] P. Sermpezis and T. Spyropoulos, "Information diffusion in heterogeneous networks: The configuration model approach," in *Proc. IEEE NetSciCom (INFOCOM Workshops)*, 2013.

[11] E. Gregori, A. Improta, L. Lenzini, and C. Orsini, "The impact of ixps on the as-level topology structure of the internet," *Elsevier Computer Communications*, vol. 34, no. 1, pp. 68–82, 2011.